EXHIBIT 2

EXHIBIT E

AKAMAI'S PRELIMINARY INVALIDITY CONTENTIONS
U.S. PATENT NO. 9,015,348[1]

Claims 1-3, 7, 10-16, and 18 of the U.S. Patent No. 9,015,348 ("'348 patent") are unpatentable under post-AIA 35 U.S.C. § 102(a)(1) in view of U.S. Patent Application Publication No. 2013/0166634 ("Holland")[2] as outlined in the chart below.

Asserted Claim	Disclosures in the Prior Art
1. A system for dynamically selecting from among a plurality of acceleration techniques implemented in a Content Delivery Network (CDN) using attributes associated with content requests, comprising:	The system of Claim 1 is depicted in the schematic diagram of FIG. 3, including the use of "a performance monitoring and measurement system for a CDN." Holland at ¶ 24, FIG. 3. *See also* Holland at ¶ 18, ¶ 30, ¶ 36, ¶¶ 95-97, detailed below.
a network interface that receives the content requests from a plurality of client systems for content stored in a plurality of edge servers distributed geographically throughout the CDN, wherein:	A "network interface" is shown in FIG. 6 as a "Communication Interface 616." *See* FIG. 6; *see also* FIGS. 1-3. "FIG. 3 shows an intermediary device 302, which is preferably (though without limitation) one of the machines 102 configured as a proxy server, as shown and described above with respect to FIGS. 1-2. Assume that the client 306 – in this example a web browser 306 running on an end-user machine – has made a request (e.g., an HTTP 'Get' request) for a particular web page to the proxy server 302 after having been mapped to that proxy server via a DNS system lookup. The proxy

[1] These preliminary invalidity contentions are provided as an exhibit to Akamai's "Preliminary Invalidity Contentions" document dated June 8, 2016, and should be read in concert with the preliminary statement, reservation of rights, and additional invalidity grounds set forth therein. As set forth more fully therein, Akamai does not acquiesce to any particular interpretation or claim construction of any feature or limitation of the '348 patent, nor does Akamai concede enablement of any particular feature or limitation of the '348 patent under 35 U.S.C. § 112. Akamai reserves the right to argue its non-infringement and invalidity positions in the alternative. Certain prior art references may contain technical features that may be the same as or substantially similar to features contained within the Accused Products. Akamai asserts in its preliminary non-infringement contentions that the claim limitations in the '348 patent do not cover those features. However, to the extent the claims are construed to cover such technical features, the '348 patent is invalid based on the presence of those same technical features in the prior art references. To the extent Akamai's non-infringement contentions require certain claim constructions, such claim constructions should not be read into Akamai's invalidity contentions; similarly, to the extent Akamai's invalidity contentions require certain claim constructions, such claim constructions shall not be read into Akamai's non-infringement contentions.

[2] Holland is available as prior art against all claims of the '348 patent under at least post-AIA 35 U.S.C. § 102(a)(1).

	server 302 obtains the page from the origin/source server 301…" Holland at ¶ 30. "In a known system such as that shown in FIG. 1, a distributed computer system 100 is configured as a content delivery network (CDN) and is assumed to have a set of machines 102 distributed around the Internet." Holland at ¶ 7 (describing a "distributed geographically" system]. The clients, servers, and other computer devices described herein may be implemented with conventional computer systems." ¶ Holland at 106. "FIG. 6 is a block diagram that illustrates hardware in a computer system 600 … The communication interface 616 provides a network link 618. The communication interface 616 may represent an Ethernet or other network interface card (NIC), a wireless interface … or other kind of input/output interface." Holland at ¶ 112. "Third party sites affiliated with content providers … offload delivery of content … to the distributed computer system 100 and, in particular, to the servers (… referred to as proxy servers … or sometimes as "edge" servers …)" Holland at ¶ 7
the content requests comprise a first content request for first content;	"Assume that the client 306--in this example a web browser 306 running on an end-user machine--has made a request (e.g., an HTTP `Get` request) for a particular web page to the proxy server 302 after having been mapped to that proxy server via a DNS system lookup. The proxy server 302 obtains the page from the origin/source server 301, or from an internally cached copy that was previously retrieved. … the web page (e.g., an HTML file …)" Holland at ¶ 30. *See also* FIG. 3, "1 – Request For Page". "Third party sites affiliated with content providers, such as web site 106, offload delivery of content (e.g., HTML, embedded page objects, streaming media, software downloads, and the like) to the distributed computer system 100 …" ¶7.
the first content request originates from a first client system in the plurality of client systems; and	"… the system includes an intermediary device, such as a web proxy server, that receives requests for content from clients, such as requests for a web page." Holland at ¶ 18. *See also* Holland at FIG. 1, ref. 122. "End user client machines 122 that desire such content may be directed to the distributed computer system to obtain that

	content more reliably and efficiently."
the first content request is associated with one or more attributes; and	"Information relating to the client and/or determined from the client's request. In addition to client IP address and other client identifiers, information derived from client request headers can be included in the beacon. Examples include a determination of geography based on client IP address, and a determination of client device identity or characteristics based on, e.g., the user agent." Holland at ¶ 100. *See also* "HTTP 'Get' request" Holland at ¶ 30. *See also*, "…User Agent header sent by the client 306 with the request." Holland at ¶ 52. *See also*, "The machine 200 [e.g., edge machine] shown in FIG. 2 may be configured to provide one or more extended content delivery features, preferably on a domain-specific, content-provider-specific basis, preferably using configuration files that are distributed to the content servers using a configuration system." Holland at ¶ 11. *See also*, "In one embodiment the beacon is sent on a content-provider-controlled hostname." Holland at ¶ 32.
an intermediate server that accelerates access to the content stored in the plurality of edge servers, the intermediate server comprising:	"FIG. 3 shows an intermediate device 302, which is preferably (though without limitation) one of the machines 102 configured as a proxy server, as shown and described above with respect to FIGs. 1-2." Holland at ¶ 30. "It should also be noted that the allocation of functions to particular machines is not limiting, as the functions recited herein may be combined or allocated amongst different machines in a variety of ways." Holland at ¶ 20.
a first interface[3] coupled to the network interface; a second interface configured to communicate with at least one of the plurality	FIG. 1, ref 115, DNS: "Distributed network agents 118 monitor the network as well as the server loads and provide network, traffic and load data to a DNS query handling mechanism 115, which is authoritative for content domains being managed by the CDN." Holland at ¶ 9. See also FIG. 1, ref 120, Metadata Control: "… illustrate a useful infrastructure for delivering and managing CDN server content

[3] The "first interface" is not referenced elsewhere in the claim.

of edge servers of the CDN;	control information and this and other content server control information (sometimes referred to as 'metadata') can be provisioned by the CDN service provider itself, or (via an extranet or the like) the content provider customer who operates the origin server." Holland at ¶ 11. "Using this kind of information, new control information to address the performance issue, in the form of an updated configuration file for example, can be fed back to the proxy server 300 (or to another component of the CDN such as the mapmaker in FIG. 1)." Holland at ¶ 36. *See also id.,* Holland at ¶ 20.
one or more memory devices having stored thereon:	FIG. 6, ref 606 Storage Device, ref 610 Main Memory. See also FIG. 2, ref 200 Machine, Holland at ¶ 10.
instructions for executing each of the plurality of acceleration techniques; and	"Generalizing, each function described above may be implemented as computer code, namely, as a set of computer instructions, executable in one or more processors to provide a special purpose machine." Holland at ¶ 107. Holland describes various acceleration techniques: "For example, a server may be configured to apply modifications to a given web page as it traverses the server (e.g., going from an origin/source to an end-user client) so as to reduce the number of requests the client has to make, to reduce the payload of the content, to accelerate client application processing/referring, to tailor the content for a particular client device (and its capabilities), or otherwise enhance the performance and functionality of the content. A wide variety of such treatments are known in the art and often referred to as 'front-end' web optimizations or as 'web content' optimizations." Holland at ¶ 14. "By way of example, … describes systems and methods for applying performance-enhancing modifications to web pages … A dynamic image delivery system is described in … [ ] describes systems and methods for streaming media and for executing a byte-based interpreter in a proxy server that can be used to modify content, to add rights management information and/or watermarks and the like." Holland at ¶ 15.

	"Other performance-enhancing aspects of the CDN platform relate to the ability to intelligently map end-user clients to servers, and to the ability to intelligently route and manage the transmission of content across the network. For example, the CDN may operate a cache hierarchy … A transport and routing mechanism for arbitrary data flows is described in … A system and method for delivery of content using intermediate nodes to facilitate content delivery is described in … A global hosting system that can utilize a network map is described in …" Holland at ¶ 16. "As previously described (and as noted in U.S. Publication No. 2011/0314091), a wide variety of performance-enhancing modifications can be automatically applied to a page (or other content). Specific examples of such treatments include in-lining content, resource consolidation, minification, image optimization, domain sharding, version control for cacheability, just-in-time loading, adjusting the time at which page scripts are run, device-adaption modifications, and so on. Other treatments might involve compression and/or de-duplification of content." Holland at ¶ 97.
a plurality of acceleration profiles, wherein each of the plurality of acceleration profiles specifies at least one of the plurality of acceleration techniques; and	"The configuration file may be delivered to the servers via the data transport mechanism. [ ] illustrate a useful infrastructure for delivering and managing CDN server content control information and this and other content server control (sometimes referred to as 'metadata')…" Holland at ¶ 11 "Given the ability to configure the CDN servers described above, a wide variety of content delivery features may be implemented in the CDN platform generally and by the CDN servers specifically." Holland at ¶ 14.
a processor configured to:	"As illustrated in FIG. 2, a given machine 200 in the CDN … comprises commodity hardware (e.g., an Intel processor) 202 …" Holland at ¶ 10.
access the one or more attributes associated with the first content request;	"Information relating to the client and/or determined from the client's request. In addition to client IP address and other client identifiers, information derived from client request headers can be included in the beacon. Examples include a determination of geography based on client IP address, and a determination of client device identity or characteristics based on, e.g., the user agent." Holland at ¶ 100.

	See also "HTTP 'Get' request" described in Holland at ¶ 30. *See also*, "Support for NavTiming can be inferred by the proxy server from the User Agent header sent by the client 306 with the request." Holland at ¶ 52. "It should be appreciated that in many cases, the system can be set up so that the client 306 sends the data back to the server 302, which then relays it to the back-end system 308, since the server 302 may be closer to the client 306." Holland at ¶ 34. "In addition to client IP address and other client identifiers, information derived from client request headers can be included in the beacon. Examples include a determination of geography based on client IP address, and a determination of client device identity or characteristics based on, e.g., the user agent." Holland at ¶ 100.
select one or more acceleration techniques from the plurality of acceleration techniques, wherein:	*See* Holland at ¶ 11.
the one or more acceleration techniques are selected based on the one or more attributes; and	"… a wide variety of performance-enhancing modifications can be automatically applied to a page (or other content). Specific examples of such treatments include in-lining content, resource consolidation, minification, image optimization, domain sharding, version control for cacheability, just-in-time loading, adjusting the time at which page scripts are run, device-adaption modifications, and so on. Other treatments might involve compression and/or de-duplification of content. The proxy server can embed identifiers in the beacon to indicate what content optimization treatments it applied to the page, or more generally, what features of the server and CDN (e.g., possibly corresponding to products offered by the CDN) were engaged when processing the page." Holland at ¶ 97. "The machine 200 shown in FIG. 2 may be configured to provide one or more extended content delivery features, preferably on a domain-specific, content-provider-specific basis, preferably using configuration files that are distributed to the content servers using a configuration system. A given configuration file preferably is XML-based and includes a set of content handling rules and directives that facilitate one or more advanced content handling features." Holland at ¶ 11.

	See also Holland at ¶ 14, ¶ 30, ¶ 52, and ¶ 97.
the one or more acceleration techniques modify the content;	*See id.* Holland at ¶ 97. *See also*, "For example, a server may be configured to apply modifications to a given web page … to tailor the content for a particular client device …" Holland at ¶ 14. *See also*, "The intermediary device obtains the requested content, modifies it (e.g., by applying one or more performance-enhancing optimizations), and serves the modified content to the client." Holland at ¶ 18.
use the one or more acceleration techniques to provide the content to the first client system;	*See* Holland at ¶ 97, ¶ 14, and ¶ 18.
receive metrics from the first client system, wherein the metrics are associated with a performance in providing the first content to a user of the first client system;	*See* FIG. 3 "5 – Beacon Data" from "End-User Client" 306. "3.1.3 Beacon Data. The beacon payload may vary, but for illustrative purposes, an example of a base set of information to be carried is tabulated below. All the NavTiming data is expressed as the delta in milliseconds between the domainLookupStart timestamp and some other NavTiming timestamp." Holland at ¶ 54. Table on page 5 of Holland illustrates beacon data from the end-user client. "… the device also inserts code into the content that will be executed by the client so as to cause the client to gather timing data reflecting how quickly the client gets and is able to process the content for the end-user, and to report that data back to a back-end processing system." The code further includes information identifying the modifications that the intermediary device made. This information is included in the reported data, so that the effect of the modification(s) on performance can be analyzed." Holland at ¶ 18. "2.0 Gathering Performance Data from Browsers" Holland at ¶ 37.

	"There are several ways for the inserted code 304 to access and collect performance data from the browser 306. …" Holland at ¶ 38. *See also* Holland at ¶ 32 ("In one embodiment the beacon is sent on a content-provider-controlled hostname.") *See also* Holland at ¶ 34 ("It should be appreciated that in many cases, the system can be set up so that the client 306 sends the data back to the server 302, which then relays it to the back-end system 308, since the server 302 may be closer to the client 306.")
dynamically update a process by which the one or more acceleration techniques are selected based on the metrics; and	FIG. 3, "8 – Updated Configurations". *See also*, FIG. 3,"5 – Beacon Data" received from End-User Client 306" by "Back End System 308"; "7 – Processed Beacon Data" sent from "Back End System 308" to "Map Maker and/or Metadata Control 320"; "8 – Updated Configurations" sent from "Map Maker and/or Metadata Control 320" to "Proxy 302". "4.0 Custom Beacon Fields." Holland at ¶ 95. "The utility of the system 300 can be extended by enabling the insertion of certain custom fields into the beacon code for reporting back with the beacon data, for use by the visualization system 310 or the mapping/control system 312 of the CDN. For example, the proxy server 302 can be configured to construct a variable indicating some information of interest, and the value could be inserted into the page with the JavaScript block in the <head> section, and subsequently collected along with the NavTiming or other data and sent to the back-end. The back-end 308 would recognize it as a custom field and automatically produce visualizations computed over the data on each branch of the test. The value might indicate a wide variety of information relating to the content delivery process, such as:" Holland at ¶ 96. "What web content optimizations the proxy server 302 applied to the page. As previously described …, a wide variety of performance-enhancing modifications can

	be automatically applied to a page (or other content). Specific examples of such treatments include in-lining content, resource consolidation, minification, image optimization, domain sharding, version control for cacheability, just-in-time loading, adjusting the time at which page scripts are run, device-adaption modifications, and so on. Other treatments might involve compression and/or de-duplification of content. The proxy server can embed identifiers in the beacon to indicate what content optimization treatments it applied to the page, or more generally, what features of the server and CDN (e.g., possibly corresponding to products offered by the CDN) were engaged when processing the page. The resulting data can be interpreted and visualized to indicate the effect of particular optimizations/features and combinations thereof on each page of a content provider." Holland at ¶ 97. "A control system 312 can be used to process the performance information and identify and address performance issues affecting the operation and performance of the CDN platform. The system may identify, for example, that certain web page optimizations/treatments are causing a problem for a particular content provider or particular category of client … Using this kind of information, new control information to address the performance issue, in the form of an updated configuration file, for example, can be fed back to the proxy server 300 (or other component of the CDN such as the mapmaker in FIG. 1)." Holland at ¶ 36. "The back-end processing system receives the data and processes it so that it can be viewed in a user interface (e.g., visualized in various ways to show the performance improvements associated with the performance features provided by the service provider), analyzed by CDN service provider personnel, and/or fed into mapping/configuration systems in order to tune the operation of the device and the CDN." Holland at ¶ 18. "The CDN can then provide this information to content providers (e.g., through an extranet portal), report the information for internal use, and/or act on this information directly." Holland at ¶ 30.
use the updated process to select acceleration techniques for subsequent requests	"The back-end processing system receives the data and processes it so that it can be … fed into mapping/configuration systems in order to tune the operation of the

associated with similar attributes.	device and the CDN." Holland at ¶ 18. *See also* Holland at ¶ 30, ¶ 36, ¶ 96, ¶ 97.
2. The system for dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 1, wherein the one or more attributes comprises a geographic location of the first client system.	"Information relating to the client and/or determined from the client's request. In addition to client IP address and other client identifiers, information derived from client request headers can be included in the beacon. Examples include a determination of geography based on client IP address, …" Holland at ¶ 100.
3. The system for dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 1, wherein the intermediate server is physically combined with one of the plurality of edge servers of the CDN.	"It should also be noted that the allocation of functions to particular machines is not limiting, as the functions recited herein may be combined or allocated amongst different machines in a variety of ways." Holland at ¶ 20.
7. The system for dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 1, wherein the one or more attributes are received by the intermediate server together with the first content request.	"Information relating to the client and/or determined from the client's request. In addition to client IP address and other client identifiers, information derived from client request headers can be included in the beacon. Examples include a determination of geography based on client IP address, and a determination of client device identity or characteristics based on, e.g., the user agent." Holland at ¶ 100. See also Holland at ¶ 14 and ¶ 97 describing device-specific attributes. See also Holland at ¶ 32 and ¶ 11 discussing hostname, domain-specific.
10. The system for dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 1, wherein the plurality of acceleration techniques comprises preemptively compressing media files to be transmitted to the client device.	"Other treatments might involve compression …" Holland at ¶ 97. "Typically, 'content delivery' refers to the storage, caching, or transmission of content – such as web pages, streaming media and applications – …" Holland at ¶ 6. "For streaming media, the machine typically includes one or more media servers, such as a Windows® Media Server (WMS) or Flash® 2.0 server, as required by the supported media formats." Holland at ¶ 10.

11. The system for dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 1, wherein the plurality of acceleration techniques comprises optimizing code associated with the first content.	"web content optimizations" Holland at ¶ 14, ¶ 97.
12. The system for dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 1, wherein the plurality of acceleration techniques comprises identifying portions of the content that are dynamic and caching portions of the content that are static.	"A system and method for delivery of content using intermediate nodes to facilitate content delivery is described in U.S. Patent No. 6,820,133, the content of which is incorporated herein by reference." Holland at ¶ 16 U.S. Patent No. 6,820,133 explains that, "[t]hus, the invention advantageously provides accelerated delivery of dynamically generated and non-static content, as well as static content." U.S. Patent No. 6,820,133 at 5:5-7.
13. The system for dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 1, wherein the processor is further configured to dynamically select one or more acceleration techniques from the plurality of acceleration techniques based attributes for each request received by the intermediate server.	"The machine 200 [e.g., edge machine] shown in FIG. 2 may be configured to provide one or more extended content delivery features, preferably on a domain-specific, content-provider-specific basis, preferably using configuration files that are distributed to the content servers using a configuration system." Holland at ¶ 11. *See also* Holland at ¶ 20.
14. A method of dynamically selecting from among a plurality of acceleration techniques implemented in a Content Delivery Network (CDN) using attributes associated with content requests, the method comprising:	*See* citations above for the following preamble of claim 1: "A system for dynamically selecting from among a plurality of acceleration techniques implemented in a Content Delivery Network (CDN) using attributes associated with content requests."
receiving a first content request for first content through a network interface, wherein:	*See* citations above for the following limitations of claim 1: "a network interface that receives the content requests from a plurality of client systems for content stored in a plurality of edge servers distributed geographically

	throughout the CDN", and "the content requests comprise a first content request for first content"
the network interface is configured to receive the content requests from a plurality of client systems for content stored in a plurality of edge servers distributed geographically throughout the CDN;	*See* citations above for the following limitations of claim 1: "a network interface that receives the content requests from a plurality of client systems for content stored in a plurality of edge servers distributed geographically throughout the CDN."
the first content request originates from a first client system in the plurality of client systems;	*See* citations above for the following limitations of claim 1: "the first content request originates from a first client system in the plurality of client systems."
the first content request is associated with one or more attributes;	*See* citations above for the following limitations of claim 1: "the first content request is associated with one or more attributes."
accessing, by an intermediate server, the one or more attributes associated with the first content request, wherein:	*See* citations above for the following limitations of claim 1: "access the one or more attributes associated with the first content request."
the intermediate server is configured to accelerate access to the content stored in the plurality of edge servers;	*See* citations above for the following limitations of claim 1: "an intermediate server that accelerates access to the content stored in the plurality of edge servers, the intermediate server."
selecting, by the intermediate server, one or more acceleration techniques from the plurality of acceleration techniques, wherein:	*See* citations above for the following limitations of claim 1: "select one or more acceleration techniques from the plurality of acceleration techniques."
the one or more acceleration techniques are selected based on the one or more attributes;	*See* citations above for the following limitations of claim 1: "the one or more acceleration techniques are selected based on the one or more attributes"
the one or more acceleration techniques modify the content; and	*See* citations above for the following limitations of claim 1: "the one or more acceleration techniques modify the content"

the one or more acceleration techniques form a first acceleration profile in a plurality of acceleration profiles;	*See* citations above for the following limitations of claim 1: "a plurality of acceleration profiles, wherein each of the plurality of acceleration profiles specifies at least one of the plurality of acceleration techniques."
using, by the intermediate server, the one or more acceleration techniques to provide the content to the first client system	*See* citations above for the following limitations of claim 1: "use the one or more acceleration techniques to provide the content to the first client system"
receiving, by the intermediate server, metrics from the first client system, wherein the metrics are associated with a performance in providing the first content to a user of the first client system;	"It should be appreciated that in many cases, the system can be set up so that the client 306 sends the data back to the server 302, which then relays it to the back-end system 308, since the server 302 may be closer to the client 306." ¶ Holland at 34.
dynamically updating, by the intermediate server, a process by which the one or more acceleration techniques are selected based on the metrics; and	*See* citations above for the following limitations of claim 1: "dynamically update a process by which the one or more acceleration techniques are selected based on the metrics."
using, by the intermediate server, the updated process to select acceleration techniques for subsequent requests associated with similar attributes.	*See* citations above for the following limitations of claim 1: "use the updated process to select acceleration techniques for subsequent requests associated with similar attributes."
15. The method of dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 14, wherein the one or more attributes affect a performance of the one or more acceleration techniques in providing the first content to a user of the first client system.	"Performance Measurements" Title. "The code includes information identifying the modifications the device made, and this is reported with the timing data, so that the effect on performance can be analyzed." Holland at Abstract. *See also* Holland at ¶ 11, ¶ 14, ¶ 30, ¶ 32, ¶ 52, ¶ 97, and ¶ 100.
16. The method of dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes	"HTTP 'Get' request" Holland at ¶ 30. *See also*, "…User Agent header sent by the client 306 with the request." Holland at

associated with the content requests of claim 14, wherein the one or more attributes are descriptive of the first client system or a network through which the first content request is received.	¶ 52. *See also*, "The machine 200 [e.g., edge machine] shown in FIG. 2 may be configured to provide one or more extended content delivery features, preferably on a domain-specific, content-provider-specific basis, preferably using configuration files that are distributed to the content servers using a configuration system." Holland at ¶ 11. *See also* Holland at ¶ 11, ¶ 14, ¶ 30, ¶ 32, ¶ 52, ¶ 97 and ¶1 00.
18. The method of dynamically selecting from among the plurality of acceleration techniques implemented in the CDN using the attributes associated with the content requests of claim 14, wherein the one or more attributes comprises a device type of the first client system.	"… device type such as mobile …" Holland at ¶ 57. *See also* Holland at ¶ 62, ¶ 66, ¶ 91, and ¶ 99.